

Description

PRIVATE RETRIEVAL OF DIGITAL OBJECTS

Inventors:

5

Feng Bao
Robert Deng
Peirong Feng

10 Technical Field

The present invention relates generally to secure and private communications enabling retrieval of digital objects from a computerized database.

Background Art

15

The World Wide Web (WWW) has evolved from a service focused on academic areas and offering scientific content into a medium for common users to access information of various origins. While surfing the Web, many users are not aware that a large number of organizations such as those in the marketing industry are gathering their private information. This information is supplemented when a user accesses a Web site, clicks a Web page, makes an electronic purchase, or downloads a file. From all the records and computerized analysis, the information collector can build a digital dossier about the users—what they do, where they go, what they read, what they buy, etc.

20

There has, therefore, been general recognition of the need for privacy protection on the Internet. One situation in which privacy is a large concern is when databases containing users' personal information are accessed. To illustrate, suppose there is a database that maintains groups of digital objects, and a user wishes to retrieve a subset of the digital objects. Two desirable constraints on database access are as follows:

25

- 1) the user can access the data the user wants, without disclosing to the database the specific digital objects actually desired; and
- 2) the user can not get any additional information from the database without the consent of the database.

5 The first constraint is referred to as *user privacy* and the second constraint is referred to as *database security*.

One example that illustrates these concepts is the task of providing electronic newspaper services over the Internet. A database maintains a collection of digital news articles. Assuming that a subscriber requests n articles, database security requires that the subscriber gets only n articles, while user privacy requires that the database cannot determine which n specific articles are retrieved by the subscriber.

The problem of private information retrieval was reviewed by B. Chor, O. Goldreich, E. Kushilevita, and M. Sudan, "Private Information Retrieval," *Proceedings of the 36th Annual Symposium on Foundations of Computer Science*, pp. 41-50, 1995. The authors were concerned with information-theoretical security and proposed a solution using multiple databases. However, the security of this solution relies on the assumption that the multiple databases do not communicate with each other, which is not guaranteed to be the case, and is additionally outside of the user's control and ability to independently verify.

Private information retrieval schemes using a single database were later proposed in B. Chor and N. Gilboa, "Computational Private Information Retrieval," *Proceedings of the 29th Annual ACM Symposium on Theory of Computing*, pp. 304-313, 1997, and E. Kushilevita and R. Ostrovsky, "Single-Database Computationally Private Information Retrieval," *Proceedings of the 38th Annual Symposium on Foundation of Computer Science*, 1997. These solutions are concerned with security based on computational assumption theory, and in particular the difficulty of factoring large prime numbers, as is done in the well-known RSA encryption scheme. However, the

computational costs of these solutions are prohibitively large due to their bit-by-bit processing approach. For example, the scheme in the Kushilevita and Ostrovsky reference requires a computational cost on the order of $O(N)$ multiplication modulo a 1024-bit number just to retrieve 1 bit of information, where N is the number of bits of data maintained by the database.

The requirement of database security in the context of private information retrieval was studied in Y. Gertner, Y. Ishai, E. Kushilevita and T. Malkin, "Protecting Data Privacy in Private Information Retrieval Schemes," *Proceedings of the 30th ACM Annual Symposium on Theory of Computing*, 1998.

All of the proposed solutions to the problem of private information retrieval described above employ the bit-by-bit processing approach. Therefore, they have only theoretical values, and are not feasible in practical applications, because of the time that would be required to solve each problem.

Therefore, what is needed is a way of allowing a user to achieve information retrieval from a database in an efficient manner while maintaining privacy.

DISCLOSURE OF INVENTION

In accordance with the present invention, there is provided a way to allow a user (102) to achieve private information retrieval from a database (104) in an efficient manner. The database (104) maintains one or more groups (106) of digital objects (202) available for users to access. A user (102) can retrieve a subset of digital objects (202) from a group (106) of digital objects (202) in the database (104) such that:

- 1) the user can access the data (202) the user (102) wants, without disclosing to the database (104) the specific digital objects (202) actually desired; and
- 2) the user (102) can not access additional information (202) from the database (104) without the consent of the database (104).

Objects (202) in the database (104) are stored in one or more different groups (106). The user (102) identifies some particular objects (202) of interest in the database (104), and additionally to which groups (106) those objects (202) belong. The user (102) then sends (302) a request to the database (104), specifying only the groups (106) containing the desired objects (202), but does not specifically identify the particular digital objects (202) desired. At this point, an electronic commerce transaction might take place, where the user (102) pays for access to a specified number of digital objects (202). The database (104) then encrypts (304) all digital objects (202) in each requested group (106) into ciphertext (206). In addition, a key (204) for each ciphertext (206) is encrypted (306). The database (104) then sends back (308) to the user (102) both the ciphertexts (206) and the associated encrypted keys (208).

At this point, the database (104) knows only that the user (102) desires one or more digital objects (202) from a particular group (106) of digital objects in the database (104), but is unable to determine which particular objects (202) are of interest.

15 The user identifies (310) the ciphertexts (206) of the desired digital objects (202), and their associated keys (208). Next, the user re-encrypts (312) the identified keys (208), and returns (314) the re-encrypted keys (506) to the database (104). The database decrypts (316) the keys (506) to the extent that it is able—i.e., the database (104) reverses the encryption it previously applied to those keys (506). However, the database (104) is
20 unable to identify which digital objects (202) the keys (506) are associated with, because the keys (512) remain encrypted with the user's encryption scheme. The database (104) now sends (318) the keys (512) back to the user (102).

Once the user (102) receives the keys (512) back from the database (104), the next step is simply to decrypt (320) them using the user's own decryption scheme (604), thus
25 revealing the unencrypted keys (204). Finally, the user (102) uses those keys (204) to decrypt (322) the appropriate digital object ciphertexts (206).

Since the database (104) is unable to determine which keys (204) it has decrypted, user (102) privacy is maintained. And, since the user (102) cannot gain access to any key (204) unless the database (104) first decrypts it, the user (102) will not be able to access any more objects (202) than are authorized. Thus, both constraints discussed
5 above have been satisfied.

The present invention does not require multiple databases. Processing is digital object (202) oriented instead of bit oriented. User (102) privacy is guaranteed without any computational constraint and without additional constraints on the "honesty" of the database (104). This means that the user's interest in specific digital objects (202) is not disclosed. The security of the database (104) is based on the assumption of the intractability of computing discrete logarithms, which forms the basis of many existing digital signature schemes and the Diffie-Hellman key exchange protocol. See W. Diffie and M. Hellman, "New directions in cryptography," *IEEE Transactions on Information Theory*, Vol. IT-22, No. 6, pp. 644-654, November 1976.

15 The present invention also provides a balance between user (102) privacy and communication cost. Communication cost can be reduced by decreasing the size of a digital object group (106), while a large digital object group (106) size gives better user (102) privacy protection.

BRIEF DESCRIPTION OF THE DRAWINGS

20 These and other more detailed and specific objects and features of the present invention are more fully disclosed in the following specification, reference being had to the accompanying drawings, in which:

Fig. 1 is a block diagram of a data access system between a (102) user and a database (104).

Fig. 2 is a block diagram of digital objects (202) inside a group (106), and corresponding keys (204), ciphertexts (206), and key ciphertexts (208) associated with the digital objects (202).

Fig. 3 is a flowchart of the operation of the illustrative embodiment of the present invention.

Fig. 4 is a block diagram illustrating the encryption of digital objects (202) into ciphertext (206), and of keys (204) into key ciphertexts (208).

Figs. 5a and 5b are block diagrams illustrating, respectively, the re-encryption of a ciphertext key (208), and the partial decryption of such a key (208).

Figs. 6a and 6b are block diagrams illustrating, respectively, the decryption of a key (512), and the decryption of a digital object ciphertext (206) using a key (204).

Fig. 7 is a block diagram of an apparatus that is a preferred embodiment of the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

A cryptographic system, or cryptosystem, has an encryption key to convert plaintext into ciphertext and a decryption key to recover the plaintext from ciphertext. If the encryption key and the decryption key are identical, the cryptosystem is called a symmetric key cryptosystem. If the encryption key and the decryption key are different and it is computationally infeasible to determine the decryption key from the mathematically-related encryption key, the cryptosystem is called an asymmetric key cryptosystem, or a public key cryptosystem. For illustrative purposes, the preferred embodiments described here make reference to symmetric key cryptosystems for encryption and decryption. It will be apparent to those skilled in the art, however, that asymmetric key cryptosystems could also be used. See, for example, A. Menezes, P. Oorschot, and S. Vanstone, *Handbook of Applied Cryptography*, CRC Press, 1996, or C.

Kaufman, R. Perlman, and M. Speciner, *Network Security - Private Communication in A Public World*, PTR Prentice Hall, Englewoor Cliffs, NJ, 1995.

For purposes of clarity, we use $e(k, m)$ to denote encryption of a digital object m with key k in a symmetric key cryptosystem; and $d(k, c)$ to denote the decryption of a
5 ciphertext c with key k in a symmetric key cryptosystem.

Fig. 1 is a model of a data access system between a user 102 and a database 104. The system contains a user 102, and a database 104. The database 104 maintains groups 106 of digital objects m 202. The user 102 wishes to access digital objects 202 in the database 104 by subscribing to the database's service, or by paying the database 104
10 with electronic cash, or by other means as required by the database 104. A connection 108 between the user 102 and the database 104 could be any standard communication media, such as the Internet or other wide area network. Further, the database 104 maintains one or more groups 106 of digital objects m , and the user 102 is interested in retrieving digital objects 202 from the group of N digital objects $\{m, i = 1, 2, \dots, N\}$ 106 in
15 the database 104. In Fig. 1, the illustrated database 104 contains Groups A through G; however it will be appreciated that the present invention is applicable to a database 104 containing any number of groups 106. It should also be noted that the particular manner in which the user 102 discovers the desired group 106 is not material to the present invention. All that is required is that the user 102, either directly or through the
20 use of client software operated by the user 102, be aware of the digital object 202 the user wants, and the group 106 in which that object 202 is located.

Fig. 2 is a block diagram of a group 106 of digital objects 202 contained within the database 104. A group 106 contains one or more digital objects 202. The number of digital objects 202 in a group 106 is determined by the operator/maintainer of the
25 database 104, and may be determined by factors not within the scope of the present invention. For purposes of the present invention, however, it will be noted from the description that decreasing the size of a group 106 reduces communication cost, but also

decreases privacy protection for the user 102. Initially, encryption is performed upon all objects 202 in the group 106, as indicated below. Thus, each digital object 202 in the group 106 will have a ciphertext 206 and a key 204, and each key 204 will additionally have an associated ciphertext 208.

5 Fig. 3 shows a flowchart of the operation of a preferred embodiment of the present invention. The database 104 and user 102 have agreed on some prime number p , such that $p = vq + 1$, where q is a large prime number, for example 160 bits in length, and v is a large integer, for example 800 bits in length. The prime number q is chosen
 10 such that p will be prime as well. When the user 102 wants to retrieve digital objects 202 from the group 106, the user 102 sends 302 a request and optionally the corresponding payment to the database 104. Upon receipt of the request, the database 104 generates 303 a random number R , $0 < R < p - 1$, and N keys k_i , $i = 1, 2, \dots, N$, for a symmetric key cryptosystem in a fashion well known in the art. One key k is associated
 15 with each digital object m . The database then encrypts 304 each digital object m_i 202 in the group 106 with k_i 204 using the symmetric key cryptosystem to obtain ciphertext $c_i = e(k_i, m_i)$, $i = 1, 2, \dots, N$ 206. Finally, the database 104 encrypts 306 the keys 204 themselves to obtain $s_i = k_i^R \bmod p$, $i = 1, 2, \dots, N$ 208.

The database 104 next transmits 308 the encrypted objects 206 and keys 208 (c_i , s_i), $i = 1, 2, \dots, N$ to the user 102. Assuming that the user 102 intends to retrieve n , $n < N$,
 20 digital objects $m_{i_1}, m_{i_2}, \dots, m_{i_n}$ 202 from the group 106, the user 102 identifies 310 the objects 206 and keys 208 desired, and generates 311 n random numbers w_j , $0 < w_j < p - 1$, and then obtains 312 n re-encrypted keys $W_j = s_{i_j}^{w_j} \bmod p$, $j = 1, 2, \dots, n$. The user 102 sends 314 W_j , $j = 1, 2, \dots, n$ and optionally the required payment to the database 104. The database 104 computes 316 and sends 318 $U_j = W_j^{1/r \bmod (p-1)} \bmod p$, $j = 1, 2, \dots, n$, back
 25 to the user 102.

The user 102 computes $k_{ij} = U_j^{1/w_j \bmod (p-1)} \bmod p$, $j = 1, 2, \dots, n$, and then decrypts c_{ij} with k_{ij} using the symmetric key cryptosystem to recover digital objects $m_{ij} = d(k_{ij}, c_{ij})$, $j = 1, 2, \dots, n$ 202.

Fig. 4 is a block diagram that further illustrates the encryption performed on a digital object 202 by the database 104. The digital object 202 and its associated key 204 are provided to the cryptosystem 406, to produce the ciphertext 206, $e(k, m)$. Similarly, using a prime number p 404, the key 204 and random number R 402, the key 204 itself is encrypted into ciphertext 208 via the cryptosystem 406.

Fig. 5a is a block diagram illustrating the process carried out by the user 102 of re-encrypting 312 the key 204. In addition to the key ciphertext 208, a prime number p 404 and random number w 502 are processed through the encryption algorithm ($s_{ij}^w \bmod p$, as described above) 504 to obtain the re-encrypted key 506.

Similarly, Fig. 5b illustrates the partial decryption 314 performed by the database 104 on the re-encrypted key 506. Using the previously-generated random number R 402 and prime number p 404, the re-encrypted key 506 is then decrypted 314 using the decryption algorithm ($W_j^{1/r \bmod (p-1)} \bmod p$) 508 to obtain the partially decrypted key U 510.

Fig. 6a illustrates the step of transforming the partially decrypted key U 510 into the unencrypted key K 204. The partially decrypted key U 510, the random number w 502, and prime number p 404 are input into the user decryption algorithm ($U_j^{1/w_j \bmod (p-1)} \bmod p$) 602, thus revealing the unencrypted key K 204.

Then, as shown in Fig. 6b, key k 204 and ciphertext c 206 are input into the cryptosystem decryption algorithm ($d(k_{ij}, c_{ij})$) 604 to obtain the digital object m 202.

Fig. 7 is a block diagram of an apparatus that is a preferred embodiment of the present invention. Note that the apparatus can be implemented either as hardware, firmware, or software. The user 102 has a user bus 726 through which each of the user

modules communicate. Similarly, the database 104 has a database bus 728. The user bus 726 and database bus 728 communicate via connection 108. The user 102 requests a group from the database 104 using the requesting module 714. The user generates random numbers using the random number generating module 718. Transmissions from the database 104 to the user 102 are received by the receiving module 716. Data is sent from the user 102 to the database 104 via the transmitting module 722. User 102 encryption is performed by the encrypting module 720, and user 102 decryption by the decryption module 724.

Focusing on the database 104 modules illustrated in Fig. 7, the database 104 generates random numbers using the random number generating module 702. Transmissions from the user 102 to the database 104 are received by the receiving module 710. Transmissions from the database 104 to the user 102 are sent by the transmitting module 708. The database 104 also has a key generating module 704 for generating keys 204, an encrypting module 706, and a decrypting module 712.

Security Considerations:

First, it can be easily seen from this description that the user 102 can obtain the desired digital objects m_j 202 by decrypting ciphertexts c_j 206 with computes $k_j = U_j^{1/w_j} \bmod (p-1) \bmod p, j = 1, 2, \dots, n$. That is, if both the database 104 and user 102 follow the protocol, the user 102 gets the desired information. However, under no circumstances is the database 104 able to pinpoint which digital objects 202 are being retrieved by the user 102. In order for the database 104 to find out which digital object 202 the user 102 is interested in retrieving, the database 104 would need to figure out which s_j 208 is being used to compute $W_j = s_j^{w_j} \bmod p$ 506 by the user 102. However, the only information available to the database 104 is $W_j = s_j^{w_j} \bmod p, j = 1, 2, \dots, n$ and $s_i, i = 1, 2, \dots, N$. Since w_j 's are randomly chosen and kept secret by the user 102, it is equally likely that all s_j 's 208 are being used in computing $W_j = s_j^{w_j} \bmod p, j = 1, 2, \dots, n$. Therefore, the user's privacy is satisfied without having to rely on any computational assumptions.

Next, we consider database 104 security. Without loss of generality, assume that the user 102 has paid and retrieved m_1, m_2, \dots, m_j 202. The user 102 then tries to recover m_{j+1} , which the user 102 is not authorized to access, without the database's 104 help. This problem is equivalent to, given s_1 208(1), k_1 204(1), s_2 208(2), k_2 204(2), ..., s_j, k_j , and s_{j+1} , finding k_{j+1} such that $s_{j+1} = k_{j+1}^R \bmod p$. One approach to solving this problem is to find R 402 from, for example, $s_j = k_j^R \bmod p$ and then compute $k_{j+1} = s_{j+1}^{1/R(p-1)} \bmod p$. But this is equivalent to solving the discrete logarithm problem, and is therefore not feasible. The second approach is to express s_{j+1} in terms of multiplication or division of s_1, s_2, \dots, s_j . Then k_{j+1} can be found from a corresponding expression in terms of k_1, k_2, \dots, k_j . However, since k_1, k_2, \dots, k_j and k_{j+1} are all independently and randomly chosen, finding the relationship between the s_j 's is also not computationally feasible.

Finally, digital objects 202 are encrypted with a symmetric key cryptosystem and the encryption keys 204 are protected using large exponentiations. To recover the digital objects 202 from the ciphertexts 206, an eavesdropper must be able to break the symmetric key cryptosystem or solve the discrete logarithm problem. Both are computationally infeasible for well-designed ciphers and exponentiations with large prime modulus.

The above description is included to illustrate the operation of the preferred embodiments and is not meant to limit the scope of the invention. The scope of the invention is to be limited only by the following claims. From the above discussion, many variations will be apparent to one skilled in the art that would yet be encompassed by the spirit and scope of the present invention.

What is claimed is: